UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/782,314	02/19/2004	Michael R. Smith	CIS0211US	7633

33031   7590   06/22/2009
CAMPBELL STEPHENSON LLP
11401 CENTURY OAKS TERRACE
BLDG. H, SUITE 250
AUSTIN, TX 78758

EXAMINER
JAKOVAC, RYAN J

ART UNIT	PAPER NUMBER
2445	

MAIL DATE	DELIVERY MODE
06/22/2009	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

<table>
<tr><td rowspan="2">Office Action Summary</td><td>Application No.</td><td>Applicant(s)</td><td></td></tr>
<tr><td>10/782,314</td><td>SMITH ET AL.</td><td></td></tr>
</table>

Office Action Summary	**Examiner**	**Art Unit**	
	RYAN J. JAKOVAC	2445	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE <u>3</u> MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.
- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133).
  Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *15 April 2009*.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-3,5-17,19-27 and 38-67* is/are pending in the application.

    4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-3,5-17,19-27 and 38-67* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☒ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

    Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

    Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

    a)☐ All   b)☐ Some * c)☐ None of:

      1.☐ Certified copies of the priority documents have been received.

      2.☐ Certified copies of the priority documents have been received in Application No. _____.

      3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

    * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1)☐ Notice of References Cited (PTO-892)

2)☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)

3)☒ Information Disclosure Statement(s) (PTO/SB/08)
   Paper No(s)/Mail Date *04/16/2009*.

4)☐ Interview Summary (PTO-413)
   Paper No(s)/Mail Date. _____.

5)☐ Notice of Informal Patent Application

6)☐ Other: _____.

DETAILED ACTION

Response to Arguments

1.   Applicant has amended to include limitations directed towards providing redundant

connections between the networked elements of the claims. However, providing redundant

connections between networked elements is a well known technique in the art used to provide a

reliable network. Providing redundant connections between networked elements as recited in the

Applicant's amended claims would have been obvious to one of ordinary skill in the art at the

time of the invention, and does not patentably distinguish the Applicant's invention from the

prior art. This further evidenced by at least paragraphs [0004-0005] if the Applicant's

specification entitled "Description of Related Art", which recites "In order to provide increased

network reliability, redundant switches and links are often included in a network", and

"…having redundant links and/or switching can provide a more reliable network."

2.   Applicant has amended claims 8, 19, 41, and 51 to include limitations directed towards

detecting whether a packet was received via a virtual network device link. At least fig. 4 of Beck

discloses that when a packet is received a determination is made whether the packet was sent to

the cluster (i.e. via the virtual network device links of fig. 2).

3.   Applicant's argues that the combination of Beck and TCP/IP does not disclose that the

first network device appends a header that identifies a port that received the packet. The

Examiner respectfully disagrees. Beck discloses transferring packets via TCP/IP. Fig. 3 shows

the TCP/IP packet and the information contained in it. Beck further discloses that data packets

are transferred between processor nodes (Beck, [0035].). The processor nodes listen for

incoming connections on certain ports (Beck, abstract, "the cluster is provided with a skinny

stack application for selecting a processor node, to which a connection will be established, **after consideration has been given to the TCP port numbers that the processor node is listening for**".). TCP/IP discloses on pg. 3 that "Each TCP segment contains the source and destination port number to identify the sending and receiving application. These two values, along with the source and destination IP addresses in the IP header, uniquely identify each connection." Given paragraph [0007] of Beck:

> "Each processor node within the cluster has the ability to distribute received data packets to an
> appropriate processor node for servicing. The processor node receiving the data packet from the
> network will hereinafter be referred to as the "receiving processor node" for that transaction.
> When a data packet arrives at the cluster, the receiving processor node first determines the type of
> the data packet. For example, most data packets correspond to the TCP/IP or UDP network
> protocols. The receiving processor node further determines whether the data packet is associated
> with an existing connection to an application running on one of the processor nodes within the
> cluster or whether a new connection should be established."

It would have been obvious to one of ordinary skill in the art at the time of the invention to combine the appending a header that identifies a port that received the packet as taught by TCP/IP with Beck *at least* to uniquely identify each connection between sending and receiving applications (i.e. between the receiving and sending processor nodes of Beck). It would have been further obvious to combine the explicit methods of TCP/IP with Beck since Beck discloses using TCP/IP protocol and this protocol is a networking standard. TCP/IP has been in existence for at least a decade prior to the Applicant's filing date of Feb. 19, 2004 as evidenced by the publication date of TCP/IP Illustrated, which is December 31, 1993.

4.   Applicant argues in summary that, regarding claim 38, the cited portions of Beck fails to

disclose a network device sending packets via one link and other packets via a second link,

where the packets are associated with identifiers identifying the same virtual link bundle.

However, Beck discloses in fig. 2 a plurality of links used for sending packets between the

processor nodes. As described above, these packets indicate the source and destination (i.e.

identifying the virtual link bundle). See fig. 3 of Beck and the discussion above.

5.   The differences between the subject matter sought to be patented and the prior art are

such that the subject matter as a whole would have been obvious at the time the invention was

made to a person having ordinary skill in the art to which said subject matter pertains.

Specification

6.   The specification is objected to as failing to provide proper antecedent basis for the

claimed subject matter. See 37 CFR 1.75(d)(1) and MPEP § 608.01(o). Correction of the

following is required: Claims 8, 19, 41, 51, and 61 recite "drop the packet from a packet flow",

however dropping the packet is not disclosed in the specification.

7.   The amendment filed 04/15/2009 is objected to under 35 U.S.C. 132(a) because it

introduces new matter into the disclosure. 35 U.S.C. 132(a) states that no amendment shall

introduce new matter into the disclosure of the invention. The added material which is not

supported by the original disclosure is as follows: Claims 8, 19, 41, 51, and 61 recite "drop the

packet from a packet flow". Applicant is required to cancel the new matter in the reply to this

Office Action.

Claim Rejections - 35 USC § 103

8.   The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all

obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in
>
> section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are
>
> such that the subject matter as a whole would have been obvious at the time the invention was made to a person
>
> having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the
>
> manner in which the invention was made.

9.   Claims 1-3, 5-17, 19-20, 22-27, 38, 40-48, 50-58, and 60-67 are rejected under 35 U.S.C.

103(a) as being unpatentable over U.S. 2001/0014097 to Beck et al (hereinafter Beck) in view of

TCP/IP Illustrated, Volume 1: The Protocols (hereinafter TCP/IP).

       Regarding claim 1, Beck teaches a system comprising:

       a virtual link bundle comprising a plurality of communication links,

       wherein a first end of each of the communication links is configured to be coupled to a

first network device (Beck, Fig. 2, Processor node A is a component of the cluster connected to

subnet S1. The communications links including the links between the processor nodes of the

cluster and the network subnet are shown in Fig. 2.);

       a second end a first one of the communication links is configured to be coupled to a first

virtual network device sub-unit within a virtual network device (Beck, Fig. 2 discloses Processor

node B (i.e. first virtual network device sub-unit) which is a node inside the cluster (i.e. virtual

network device). The communications links are show connected via the subnet s1 between the

network router and the processor nodes.);

a second end of a second one of the communication links is configured to be coupled to a

second virtual network device sub-unit within the virtual network device (Beck, Fig. 2 discloses

Processor node C (i.e. second virtual network device sub-unit) which is a node inside the cluster

(i.e. virtual network device). The communications links are shown connected via the subnet s1

between the network router and the processor nodes.);

the first one of the communication links and the second one of the communications links

provide redundant connections between the first network device and the first virtual network

device (Beck, Fig. 2, multiple communications links between cluster and processor nodes. See

also, [0026], multiple subnets couple the processing nodes. Further, providing redundant

connections between networked elements is a well known technique in the art used to provide a

reliable network. Providing redundant connections between networked elements would have

been obvious to one of ordinary skill in the art at the time of the invention);

the first network device comprises a plurality of ports; each of the ports is configured to

communicate packets with a respective client (Beck, abstract, [0005], [0011].)

the first network device is configured to append a header to a packet before sending the

packet to the virtual network device via one of the communication links (Beck, [0009], a

receiving processor node modifies the header of a packet being sent to another processing node.).

TCP/IP discloses the header identifies one of the ports having received the packet

(TCP/IP, pg. 2-4, the format of the TCP header includes a source port number (i.e. identifies the

receiving port).)

Therefore it would have been obvious to one of ordinary skill in the art at the time of

invention to combine information in a header which identifies one of the receiving ports as

taught by TCP/IP with the system of Beck in order to identify the source port and would have

been further obvious since including this information in a header is included in standard TCP

formatting which comprises well known techniques in the art.

Regarding claim 2, the combination of Beck and TCP/IP teaches the system of claim 1,

further comprising the first network device, wherein the first network device is configured to

select a communication link of the plurality of communication links on which to send a

particular packet (Beck, Paragraph [0009], A data packet is sent to the network for delivery to a

particular node (i.e. along a particular link to the node)).

Regarding claim 3, The combination of Beck and TCP/IP teaches the system of claim 2,

wherein each packet sent between the virtual network device and the first network device is sent

via only a one of the communication links (Beck, Paragraph [0030], A destination address is

specified for a particular node, thus packets are sent via the specific link associated with an

individual processor.).

Regarding claim 5, The combination of Beck and TCP/IP teaches the system of claim 1,

further comprising the first virtual network device sub-unit, wherein the first virtual network

device sub-unit is configured to identify whether a one of the communication links is coupled to

another virtual network device sub-unit within the virtual network device (Beck, Paragraph

[0011], When a link has failed, the address of the processor node is acquired by another

processor node for the duration of the failure.).

Regarding claim 6, The combination of Beck and TCP/IP teaches the system of claim 1,

further comprising the first virtual network device sub-unit (Beck, Fig. 2, Processor node B) and

the second virtual network device sub-unit (Beck, Fig. 2, processor node C), wherein the first

virtual network device sub-unit and the second virtual network device sub-unit are configured to

communicate packets with each other via a virtual network device link (Beck, Paragraph [0027],

Processor nodes distribute packets within the cluster.).

Regarding claim 7, The combination of Beck and TCP/IP teaches the system of claim 1,

wherein the communication links are configured to be managed as a single link (Beck, Paragraph

[0004], Cluster alias addresses are used to make the cluster appear to be a single node.).

Regarding claim 8, The combination of Beck and TCP/IP teaches a system comprising: a

first virtual network device sub-unit (Beck, Fig. 2, Processor Node A) comprising: a first

interface (Beck, Fig. 2, Processor interface 20a.); and a controller coupled to the first interface

and configured to forward packets received via the first interface (Beck, Paragraph [0027], The

processor node distributes packets to other nodes within the cluster.), wherein the first interface

is identified by a first logical identifier, a second interface is identified by the first logical

identifier (Beck, Fig. 2, Interfaces are associated with an IP address.), an interface bundle

comprises the first interface and the second interface (Beck, Fig. 2 discloses a grouping of

processor nodes called a cluster with interconnected communication links.), and the second

interface is comprised in a second virtual network device sub-unit (see also paragraph Beck,

[0027]),

The controller is configured to detect whether a packet was received via a virtual network

device link (Beck, fig. 4, packet is received and a determination is made whether it was sent to

the cluster (i.e. via the virtual network device links of fig. 2.).), and the first interface is

configured to drop the packet from a packet flow being sent via the first interface if the packet

was received via the virtual network device link (Beck, [0039-0041], incoming packets are

received and determination is made as to which processor node should handle the incoming

connection. Connections are routed to the appropriate processor node (i.e. the packet flow is

filtered).).

Regarding claim 9, The combination of Beck and TCP/IP teaches the system of claim 8,

further comprising the second virtual network device sub-unit (Beck, Fig. 2 discloses a system

including multiple processor nodes (i.e. second virtual network device sub-unit.).

Regarding claim 10, The combination of Beck and TCP/IP teaches the system of claim 9,

wherein the first virtual network device sub-unit is configured to maintain consistent forwarding

information with the second virtual network device sub-unit (Beck, Paragraph [0027], The

processor node distributes packets to other nodes within the cluster.).

Regarding claim 11, The combination of Beck and TCP/IP teaches the system of claim

10, wherein the controller is configured to perform control protocol processing for the first

interface according to a routing protocol running on the interface bundle (Beck, Paragraph

[0076], Nodes in the cluster establish databases containing the network layer addresses used by

each processor node in the cluster.), the controller is configured to provide information generated

when performing the control protocol processing to a secondary controller comprised in the

second virtual network device sub-unit (Beck, Paragraph [0053-0054], The Mbuf chain data

structure is sent across a cluster interconnect to a processor node.), and the secondary controller

is configured to use the information to manage the second interface (Beck, Paragraph [0055],

The Mbuf chain is queued for service of packets on the second node.).

Regarding claim 12, The combination of Beck and TCP/IP teaches the system of claim 8,

wherein the controller is configured to lookup a destination address of a first packet in a lookup

table (Beck, Paragraph [0040-0041], The receiver listens on a destination port number. The

processor node looks up receiver applications in a lookup table.), and if the lookup table returns

the first logical identifier, the first virtual network device sub-unit is configured to prioritize

sending the first packet via the first interface over sending the first packet via the second

interface (Beck, Paragraph [0041-0044], A database is maintained from which a node is selected

based on selection weights which include IP address (i.e. logical identifier).).

Regarding claim 13, The combination of Beck and TCP/IP teaches the system of claim

12, wherein if the lookup table returns the first logical identifier, the first virtual network device

sub-unit is configured to send the first packet via the first interface instead of sending the packet via the second interface, unless one or more of the first interface and a link coupled to the first interface are failed (Beck, Paragraph [0076], Nodes in the cluster establish databases containing the network layer addresses used by each processor node in the cluster. If the processor node crashes, another node takes over.).

Regarding claim 14, The combination of Beck and TCP/IP teaches the system of claim 13, wherein the first virtual network device sub-unit comprises a plurality of interfaces, more than one of the interfaces are each comprised in the interface bundle, and the more than one of the interfaces comprises the first interface (Beck, Paragraph [0062], Each node may contain more than one interface. Fig. 7 discloses nodes having multiple interfaces. Nodes with multiple interfaces are comprised in the cluster (i.e. interface bundle).).

Regarding claim 15, The combination of Beck and TCP/IP teaches the system of claim 14, wherein if each respective communication link coupled to the more than one of the interfaces fails, the first virtual network device sub-unit is configured to forward the first packet via the second interface comprised in the second virtual network device sub-unit (Beck, Paragraph [0076], Nodes in the cluster establish databases containing the network layer addresses used by each processor node in the cluster. If the processor node crashes, another node takes over.).

Regarding claim 16, The combination of Beck and TCP/IP teaches the system of claim 8, wherein the first virtual network device sub-unit (Beck, Fig. 2, Processor Node A) is coupled to

the second virtual network device sub-unit by a virtual network device link (Beck, Fig. 2 discloses the network links between the processors in the cluster including the network subnet.).

Regarding claim 17, The combination of Beck and TCP/IP teaches the system of claim 16, wherein the first virtual network device sub-unit is configured to learn that a source address of the second packet is behind the first interface, in response to receiving a second packet via the virtual network device link (Beck, Paragraph [0009], The packet is sent from one node to another over the network.).

Regarding claim 19, The combination of Beck and TCP/IP teaches a system comprising: a virtual link bundle (Beck, Fig. 2 discloses subnet S1 connected to processor nodes of a cluster. The subnet and the connections to the processor nodes comprise the virtual link bundle.); a first virtual network device sub-unit (Beck, Fig. 2, Processor node B), wherein the first virtual network device sub-unit is configured to detect whether a packet was received via a virtual network device link (Beck, fig. 4, packet is received and a determination is made whether it was sent to the cluster (i.e. via the virtual network device links of fig. 2.).); and a second virtual network device sub-unit (Beck, Fig. 2, Processor node C), wherein a first interface of the first virtual network device sub-unit is coupled to the virtual link bundle (Beck, Fig. 2, Processor node B has an processor interface 20b which is connected to the subnet.), a second interface of the second virtual network device sub-unit is coupled to the virtual link bundle (Beck, Fig. 2, Processor node C has an processor interface 20C which is connected to the subnet.), and each of the first interface and the second interface is identified by a first logical identifier (Beck, Fig. 2,

The interfaces are identified by IP address, specified in the figure as S1.B and S1.C.), and the
first interface is configured to drop the packet from a packet flow being sent via the first
interface if the second packet was received via the virtual network device link (Beck, [0039-
0041], incoming packets are received and determination is made as to which processor node
should handle the incoming connection. Connections are routed to the appropriate processor
node (i.e. the packet flow is filtered).).

Regarding claim 20, The combination of Beck and TCP/IP teaches the system of claim 19,
further comprising: a network device coupled to the first virtual network device sub-unit and the
second virtual network device sub-unit by the virtual link bundle (Beck, Fig. 2 discloses a node
(i.e. network device) connected to the processor nodes (i.e. first and second virtual network
devices.).

Regarding claim 22, The combination of Beck and TCP/IP teaches the system of claim
19, wherein a primary controller comprised in the first virtual network device sub-unit is
configured to perform control protocol processing for the first interface according to a routing
protocol running on the virtual link bundle (Beck, Paragraph [0007], Nodes distribute TCP
packets to other processing nodes for servicing.), the primary controller is configured to send
information generated by performing the control protocol processing to a secondary controller
comprised in the second virtual network device sub-unit, and the secondary controller is
configured to use the information to manage the second interface (Beck, Paragraph [0055], The
Mbuf chain is sent to a node from a second node and queued for service on the second node.).

Regarding claim 23, The combination of Beck and TCP/IP teaches the system of claim 19, wherein the first virtual network device sub-unit is configured to lookup a destination address of a packet in a lookup table (Beck, Paragraph [0041], The node accesses a lookup table in accordance with which port and a node (i.e. destination address) is listening on for incoming packets.), and if the lookup table returns the first logical identifier, the first virtual network device sub-unit is configured to prioritize sending the packet via the first interface over sending the packet via the second interface (Beck, Paragraph [0076], Nodes in the cluster establish databases containing the network layer addresses used by each processor node in the cluster. If the processor node crashes, another node takes over.).

Regarding claim 24, The combination of Beck and TCP/IP teaches the system of claim 23, wherein each of a plurality of interfaces comprised in the first virtual network device sub-unit is coupled to a respective communication link comprised in the virtual link bundle, and the interfaces comprise the first interface (Beck, Paragraph [0062], Each node may contain more than one interface. Fig. 7 discloses nodes having multiple interfaces. Nodes with multiple interfaces are comprised in the cluster (i.e. interface bundle).).

Regarding claim 25, The combination of Beck and TCP/IP teaches the system of claim 24, wherein if each respective communication link coupled to the interfaces fails, the first virtual network device sub-unit is configured to send the packet via the second interface comprised in the second virtual network device sub-unit (Beck, Paragraph [0076], Nodes in the cluster

establish databases containing the network layer addresses used by each processor node in the

cluster. If the processor node crashes, another node takes over.).

Regarding claim 26, The combination of Beck and TCP/IP teaches the system of claim

23, wherein the first virtual network device sub-unit (Beck, Fig. 2, Processor Node A) is coupled

to the second virtual network device sub-unit by a virtual network device link (Beck, Fig. 2

discloses the network links between the processors in the cluster including the network subnet.).

Regarding claims 38, 48, and 58 The combination of Beck and TCP/IP teaches a method,

system, and a computer readable medium comprising: sending a first packet via a first link of a

virtual link bundle if a destination identifier associated with the first packet identifies the virtual

link bundle (Beck, fig. 2 displays links between a plurality of processors. The processors

communicate over the communication links disclosed in fig. 2. See also [0039-0041].); and

sending a second packet via a second link of the virtual link bundle if a destination identifier

associated with the second packet identifies the virtual link bundle (Beck, Paragraph [0009],

When a receiving node determines which processor node to send to, it sends the data packet over

the network for delivery to the processor node. See also fig. 2.), wherein a single network device

performs both the sending the first packet and the sending the second packet (Beck, [0039-0041],

fig. 2, processor node A sends to both processor B and processor C.), the first link is coupled to a

first virtual network device sub-unit, and the second link is coupled to a second virtual network

device sub-unit (Beck, Fig. 2 discloses a cluster containing processor nodes B (first virtual

network sub-unit) and C (second virtual network device sub-unit) which are connected via links associated with the subnet depicted in Fig. 2.).

Regarding claims 40, 50, and 60 The combination of Beck and TCP/IP teaches the method of claim 39, the system of claim 49, and the computer readable medium of claim 59 further comprising: appending a header to the first packet (Beck, Paragraph [0009], The data packet's header is modified before it is delivered across the network to the processor node (i.e. via the first link).), wherein the header identifies which port of a plurality of ports received the first packet (Beck, Paragraph [0030], The TCP/IP header identifies the source node, the destination node, the sending port, the destination port, and the protocol being used.), and sending the first packet via the first link comprises sending the header via the first link.

Regarding claims 41, 51, and 61 The combination of Beck and TCP/IP teaches a method, system, and the computer readable medium comprising: receiving a packet, wherein a destination identifier for the packet identifies an interface bundle (Beck, Paragraph [0004], Cluster alias addresses are used to make the cluster appear to be a single node.), and the interface bundle comprises a first interface (Beck, Fig. 2 discloses a cluster comprising processor nodes, each with a processor interface. The processor interface of processor node B is the first interface.); detecting whether the packet was received via a virtual network device link (Beck, fig. 4, packet is received and a determination is made whether it was sent to the cluster (i.e. via the virtual network device links of fig. 2.).); and dropping the packet from a packet flow being sent via the first interface if the packet was received via the virtual network device link (Beck, [0039-0041],

incoming packets are received and determination is made as to which processor node should

handle the incoming connection. Connections are routed to the appropriate processor node (i.e.

the packet flow is filtered).).

      Regarding claims 42, 52, and 62 The combination of Beck and TCP/IP teaches the

method of claim 41, the system of claim 51, and the computer readable medium of claim 61

further comprising: sending the packet via the first interface if the packet was not received via

the virtual network device link (Beck, Paragraph [0058], Nodes forward packets to each other via

the cluster interconnect.).

      Regarding claims 43, 53, and 63 The combination of Beck and TCP/IP teaches the

method of claim 42, the system of claim 52, and the computer readable medium of claim 62

further comprising: maintaining consistency between a lookup table comprised in a first virtual

network device sub-unit and a second lookup table comprised in a second virtual network device

sub-unit (Beck, Paragraph [0076], Each processor node within the cluster establishes a database

containing the network layer address used by each of the processor nodes in the cluster.).

      Regarding claims 44, 54, and 64 The combination of Beck and TCP/IP teaches the

method of claim 42, the system of claim 52, and the computer readable medium of claim 62

further comprising performing control protocol processing for the interface bundle at a primary

controller comprised in a first virtual network device sub-unit (Beck, Paragraph [0007], Nodes

distribute TCP packets to other processing nodes for servicing.), wherein the first interface is

comprised in the first virtual network device sub-unit (Beck, Fig. 2 discloses a cluster comprised

of nodes (i.e. first virtual network device sub-unit) each of which has a processor interface.).

Regarding claims 45, 55, and 65 The combination of Beck and TCP/IP teaches the

method of claim 44, the system of claim 54, and the computer readable medium of claim 64

further comprising: managing a second interface of the second virtual network device sub-unit in

response to information generated by the performing the control protocol processing (Beck,

Paragraph [0055], The Mbuf chain is queued for service of packets on the second node.),

wherein the second interface is comprised in the interface bundle (Fig. 2 discloses a cluster (i.e.

interface bundle) comprising processor node C with processor interface 20c.).

Regarding claims 46, 56, and 66 The combination of Beck and TCP/IP teaches the

method of claim 45, the system of claim 55, and the computer readable medium of claim 65

further comprising: looking up a destination address of a second packet in a lookup table (Beck,

Paragraph [0030-0036], The TCP/IP header identifies the source node, the destination node, the

sending port, the destination port, and the protocol being used. TCP port numbers are used to

designate queues into which arriving packets are placed for service by nodes. Paragraph [0076],

Each processor node has a database containing the network layer addresses used by each

processor in order to take over in case of a failure.), and if the lookup table returns the

destination identifier, sending the sending packet via the first interface of the first virtual network

device sub-unit instead of sending the packet via the second interface of the second virtual

network device sub-unit (Beck, Paragraph [0076], If a processor node crashes, another takes over for it.).

Regarding claim 47, The combination of Beck and TCP/IP teaches the method of claim 41, further comprising: learning that a source address of the packet is behind a local interface, in response to receiving the packet via the virtual network device link (Beck, Paragraph [0009], The packet is sent from one node to another over the network. See also [0039-0041].).

Regarding claim 57, The combination of Beck and TCP/IP teaches the system of claim 51, further comprising: means for learning that a source address of the packet is behind a local interface, in response to receiving the packet via the virtual network device link (Beck, Paragraph [0009], The packet is sent from one node to another over the network. See also [0039-0041]).

Regarding claim 67, The combination of Beck and TCP/IP teaches the computer readable medium of claim 61, wherein the program instructions are further executable to: learn that a source address of the packet is behind a local interface, in response to detecting reception of the packet via the virtual network device link (Beck, Paragraph [0009], The packet is sent from one node to another over the network. See also [0039-0041]).

10.   Claims 21, 39, 49, and 59 are rejected under 35 U.S.C. 103(a) as being unpatentable over

the combination of Beck and TCP/IP in view of U.S. 6,735,205 to Mankude et al (hereinafter

Mankude).

        Regarding claim 21, the combination of Beck and TCP/IP teaches the system of claim 20,

Mankude teaches wherein the network device is configured to use a hash-based load-sharing

algorithm to select one of a plurality of communication links comprised in the virtual link bundle

(Mankude, Col. 7, line 10-25, In order to select a server node to forward the packet to, the

system hashes the source address of the client. The hashing selects an entry to identify a server

node within a clustered computing system.), and The combination of Beck and TCP/IP teaches

the network device is configured to send a packet via the selected one of the communication

links (Beck, Paragraph [0027], Nodes distribute packets to other nodes within the cluster.).

        Therefore it would have been obvious to one of ordinary skill in the art at the time of

invention to combine teaches wherein the network device is configured to use a hash-based load-

sharing algorithm to select one of a plurality of communication links comprised in the virtual

link bundle as taught by Mankude with the system of the combination of Beck and TCP/IP in

order to select server node to forward a packet to (Mankude, Col. 7, line 10-25).

        Regarding claims 39, 49, and 59, The combination of Beck and TCP/IP teaches the

method of claim 38, the system of claim 48, and the computer readable medium of claim 58

further comprising: selecting the first link from a plurality of links comprised in the virtual link

bundle (Beck, Paragraph [0009], When a receiving node determines which processor node to

send to, it sends the data packet over the network for delivery to the processor node.), Mankude

teaches wherein the selecting comprises performing a hash-based algorithm (Mankude, Col. 7,

line 10-25, In order to select a server node to forward the packet to, the system hashes the source

address of the client. The hashing selects an entry to identify a server node within a clustered

computing system.).

Conclusion

11.   **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time

policy as set forth in 37 CFR 1.136(a).

        A shortened statutory period for reply to this final action is set to expire THREE

MONTHS from the mailing date of this action. In the event a first reply is filed within TWO

MONTHS of the mailing date of this final action and the advisory action is not mailed until after

the end of the THREE-MONTH shortened statutory period, then the shortened statutory period

will expire on the date the advisory action is mailed, and any extension fee pursuant to 37

CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event,

however, will the statutory period for reply expire later than SIX MONTHS from the mailing

date of this final action.

        Any inquiry concerning this communication or earlier communications from the

examiner should be directed to RYAN J. JAKOVAC whose telephone number is (571)270-5003.

The examiner can normally be reached on Monday through Friday, 7:30 am to 5:00 pm EST.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's
supervisor, Vivek Srivastava can be reached on 571-272-7304. The fax phone number for the
organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent
Application Information Retrieval (PAIR) system. Status information for published applications
may be obtained from either Private PAIR or Public PAIR. Status information for unpublished
applications is available through Private PAIR only. For more information about the PAIR
system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR
system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would
like assistance from a USPTO Customer Service Representative or access to the automated
information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/RJ/

/VIVEK SRIVASTAVA/
Supervisory Patent Examiner, Art Unit 2445